arXiv:1610.09077v2 [cs.IR] 5 Nov 2016

Integrating Topic Models and Latent Factors for Recommendation

Danis J. Wilson* Wei Zhang†

Abstract

The research of personalized recommendation techniques today has mostly parted into two mainstream directions, i.e., the factorization-based approaches and topic models. Practically, they aim to benefit from the numerical ratings and textual reviews, correspondingly, which compose two major information sources in various real-world systems. However, although the two approaches are supposed to be correlated for their same goal of accurate recommendation, there still lacks a clear theoretical understanding of how their objective functions can be mathematically bridged to leverage the numerical ratings and textual reviews collectively, and why such a bridge is intuitively reasonable to match up their learning procedures for the rating prediction and top-N recommendation tasks, respectively.

In this work, we exposit with mathematical analysis that, the vector-level randomization functions to coordinate the optimization objectives of factorizational and topic models unfortunately do not exist at all, although they are usually pre-assumed and intuitively designed in the literature. Fortunately, we also point out that one can avoid the seeking of such a randomization function by optimizing a Joint Factorizational Topic (JFT) model directly. We apply our JFT model to restaurant recommendation, and study its performance in both normal and cross-city recommendation scenarios, where the latter is an extremely difficult task for its inherent cold-start nature. Experimental results on real-world datasets verified the appealing performance of our approach against previous methods, on both rating prediction and top-N recommendation tasks.

1 Introduction

The vast amount of items in various web-based applications has made it an essential task to construct reliable Personalized Recommender Systems (PRS) [1, 2]. With the ability to leverage the wisdom of crowds, Collaborative Filtering (CF)-based [3, 4] approaches have achieved significant success and wide application, especially for those Latent Factor Models (LFM) [5] based on Matrix Factorization (MF) [6, 7, 8, 9, 10, 11, 12, 13, 14] techniques, which model the preferences of users and items collectively through multivariate hidden factors learned from star ratings.

Recently, researchers have been putting attention on another important information source in many online systems, namely, the textual user reviews. Usually, the ratings and reviews come in pairs in many typical applications, *e.g.*, Amazon and Yelp. While the ratings act as integrated indicators of user attitudes towards products, the reviews serve as more detailed explanations of what aspects users care about and why the corresponding rating is made [15, 16].

As such, the application of Topic Models [17] has gained attention to leverage the textual reviews for personalized recommendation, especially the frequently used Latent Dirichlet Allocation (LDA) [18] technique and its variants, for their ability to extract latent topics from reviews that represent the actual factors users care about when making numerical ratings [19, 20, 21]. This further leads to the recent research direction to bridge the LFM and LDA models, which makes use of the ratings and reviews collectively for personalized recommendation [19, 20, 22, 23, 24, 25].

However, without a clear mathematical understanding of how the objective functions of LFM and LDA interact with each other when bridged for unified model learning, current approaches have to base themselves on pre-assumed designations to bridge their heterogenous objective functions. e.g., to transform the latent factors in LFM to topic distributions in LDA through manually designed randomization functions such as logistic normalization [19], or to assume the factors and topics be the same vector sampled from mixture Gaussian distributions [20], etc.

In this work, we study the mathematical relations between the probability of recommending an item and the estimated user-item correlations from LFM or LDA. Based on this, we prove that a multiplicatively monotonic randomization function that transforms LFM latent factors to LDA topic distributions actually does not exist at all. As a result, although some normalization-based transformations seem intuitional, they actually make the objective functions of LFM and LDA conflict with each other during optimization, where a higher value of log-likelihood in LDA component may force a

*Arcadia University. danisjwilson@gmail.com
†Carnegie Mellon University. wei@andrew.cmu.edu

lower rating prediction in LFM component, which is not favoured when bridging the two models.

Fortunately, we further find that instead of transforming a latent factor to a topic distribution separately, we can transform the product of latent factors in LFM to the corresponding product of topic distributions in LDA as a whole, thus to avoid the seeking of a theoretically nonexistent randomization function. Based on these findings, we propose the Joint Factorizational Topic (JFT) model to bridge LFM and LDA, so as to adopt the numerical ratings and textual reviews collectively, and at the same time guarantee the inner-model consistency between the LFM and LDA components.

To validate the performance, we conduct extensive experiments on (both normal and cross-city) restaurant recommendation, which is a representative task in many Location-Based Services (LBS), yet it is also extremely difficult for its inherent nature of cold-start in cross-city settings. As a result, this problem setting validates the performance of our approach and highlights the effect of leveraging the power of ratings and reviews jointly. What's more, the concept of location in this work can be further generalized to item categories, which makes it possible for our JFT framework to bridge factorizational and topical models for cross-domain recommendation.

2 Preliminaries and Definitions

2.1 Latent Factor Models (LFM) Latent Factor Models (LFM) [4] attempt to encode user and item preferences in a latent factor space so as to estimate the user-item relations for rating prediction, which account for many of the frequently used Matrix Factorization (MF) [6] techniques. Among those, a 'standard' and representative formalization [26] predicts the user-item ratings $r_{u,i}$ with user/item biases and latent factors by,

$$rate(u,i) = \alpha + \beta_u + \beta_i + \gamma_u \cdot \gamma_i \tag{2.1}$$

where α is the global offset, β_u and β_i are user and item biases, γ_u and γ_i are the K-dimensional latent factors of user u and item i, respectively, and "$\cdot$" denotes vector multiplication. Intuitively, γ_u can be interpreted as the preference of user u to some latent factors, while γ_i is the property embedding of item i on those latent factors. Based on a set of observed training records $\mathcal{R}$, the model is typically targeted with the goal of providing accurate rating predictions, where we determine the parameter set $\Theta = \{\alpha, \beta_u, \beta_i, \gamma_u, \gamma_i\}$ with the following minimization problem,

$$\Theta = \underset{\Theta}{argmin} \sum_{r_{u,i} \in \mathcal{R}} \big(rate(u,i) - r_{u,i}\big)^2 + \lambda \Omega(\Theta) \tag{2.2}$$

and $\Omega(\Theta)$ is a regularization term. A variety of methods exist to minimize Eq.(2.2), for example, Stochastic Gradient Descent (SGD) or Alternating Least Squares (ALS) [26]. However, this model merely takes into account the numerical ratings and leaves out the textual reviews, which is information-rich and may well help to provide better recommendations.

Table 1: The table of notations in this work.

$r_{u,i}$	The numerical rating given by user u to item i
U, I, C	The total number of Users, Items, or Cities
K	The dimensionality of latent factors in LFM or the number of latent topics in LDA
α	The global offset in LFM
β_u, β_i	The rating bias parameter of user u or item i
$\gamma_u, \gamma_i, \gamma_c$	The latent factors for user u, item i, or city c
$\mathcal{R}, \mathcal{D}$	The whole rating ($\mathcal{R}$) and review ($\mathcal{D}$) corpus
d_u, d_i, d_c	The set of all reviews of user u, item i, or city c
θ_d	The topic distribution of document d
ϕ_z	The word distribution of latent topic z
$z_{u,i,j}$	The corresponding topic of the jth word in the review given by user u towards item i
$w_{d,j}$	The jth word of document d

2.2 Latent Dirichlet Allocation (LDA) Different from LFM, the LDA model attempts to learn a number of K latent topics from documents (textual reviews in this work), where each word w is assigned to a topic z_w, and each topic z is associated with a word distribution ϕ_z. Based on this, each document $d \in \mathcal{D}$ is represented with a K-dimensional topic distribution θ_d, where the j-th word $w_{d,j}$ in document d discusses its corresponding topic $z_{d,j}$ with probability $\theta_{d,z_{d,j}}$. It is usually convenient to also define the word distribution $\phi_{z,w}$, which is the probability that word w is used for topic z in the whole corpus $\mathcal{D}$. The final model conducts parameter learning by maximizing the likelihood of observing the whole $\mathcal{D}$:

$$P(\mathcal{D}|\theta, \phi, z) = \prod_{d \in \mathcal{D}} \prod_{j=1}^{L_d} \theta_{d,z_{d,j}} \phi_{z_{d,j},w_{d,j}} \tag{2.3}$$

where L_d is the length (number of words) of document d. Intuitionally, we are multiplying the probability of seeing a particular topic in θ_d with the likelihood of seeing a particular word given the topic to estimate the likelihood of seeing the whole corpus.

2.3 Randomization Function Let $\gamma \in \mathbb{R}^K$ be an arbitrary vector and $\theta \in [0,1]^K$ be a stochastic vector, where their dimensions are the same K as the latent factors γ_u, γ_i and latent topics θ_d in the previous subsections. According to the definition, we have $0 \leq \theta_k \leq 1$ and $\|\theta\|_1 = \sum_{k=1}^K \theta_k = 1$. The target of a randomization function $f : \mathbb{R}^K \to \mathbb{R}^K$ is to convert

an arbitrary vector γ to a probabilistic distribution $\theta = f(\gamma)$. The inherent nature of a randomization function is the key component to bridge the gap between LFM and LDA models, which links the latent factors γ in LFM to the topic distributions θ in LDA, and thus makes it possible to model the numerical ratings and textual reviews in a joint manner.

In the background of personalized recommendation, a desired randomization function is expected to be *monotonic* in the sense that it preserves the orderings, so that the largest value of γ should also correspond to the largest value in θ, thus the dimensions of the LFM model and the LDA model are inherently aligned during the model learning process to express the user-item relations in a shared feature space. As a result, the basic properties of a randomization function $f(\cdot)$ can be summarized as follows:

$$\begin{cases} 0 \leq f(\gamma)_i \leq 1, \|f(\gamma)\|_1 = 1 \\ \gamma_i < \gamma_j \rightarrow f(\gamma)_i < f(\gamma)_j \end{cases}, \forall \gamma \in \mathbb{R}^K, 1 \leq i, j \leq K \tag{2.4}$$

For example, in [19] the authors designed a randomization function as:

$$\theta_k = f(\gamma)_k = \frac{\exp(\kappa\gamma_k)}{\sum_{k'} \exp(\kappa\gamma_{k'})} \tag{2.5}$$

which conducts logistic normalization on a latent factor. In the following, we investigate the relationship between the objective functions of LFM and LDA, and further point out the properties required on a randomization function to coordinate the models when bridging the two different functions.

3 Bridging Factors and Topics

Though various personalization strategies based on different optimization objectives (e.g., RMSE, likelihood, ranking loss, etc) may be designed for more accurate recommendation, the inherently underlying personalization indicator that we are actually considering is $P(i|u)$, which is the probability or tendency for us to recommendation an item i given a user u. Various different strategies are actually estimating this probability either explicitly or implicitly to rank the items for each user so as to construct the personalized recommendation list. In this section, we bridge the LFM and LDA models by expositing how their objective functions rank a list of items based on $P(i|u)$.

3.1 Probability of Item Recommendation

In the Latent Factor Model (LFM), a recommendation list is constructed in descending order of the predicted ratings $rate(u, i)$ for a given user, which means that an item i with a higher rating prediction on user u also gains a higher probability of being recommended $P(i|u)$. As a result,

$$P_{\text{LFM}}(i|u) \propto rate(u, i) \propto \gamma_u \cdot \gamma_i \tag{3.6}$$

where $\propto$ denotes a positive correlation, and we leave out the parameters α, β_u and β_i because they are constants given a user and an item in model learning [26].

In Latent Dirichlet Allocation (LDA) for personalized recommendation, each user or item is represented by its corresponding set of textual reviews d_u or d_i, and the underlying intuition models the topical correlation between them by estimating the potential review $d_{u,i}$ that a user may write on an item, based on the topical distributions θ_{d_u} and θ_{d_i}. To simplify the notations, we use u, i, and k to denote the user or item document representations d_u, d_i and the k-th latent topic z_k interchangeably, and we thus have:

$$P(k|u) = \theta_{u,k} \text{ and } P(k|i) = \theta_{i,k} \tag{3.7}$$

The LDA model conducts likelihood maximization on each observed review $d_{u,i}$ given the corresponding user u and item i, and the embedded topical distribution represents the probability of observing each topic k, which is:

$$P(k|u, i) = \theta_{d,k} \tag{3.8}$$

LDA applies an indirect causal effect from users to items via latent topics [18], which means that user u and item i are conditionally independent given topic k, i.e., $P(u|i, k) = P(u|k)$, and this further gives us:

$$P(u, i, k) = \frac{P(u, k)P(i, k)}{P(k)} \tag{3.9}$$

By applying Eq.(3.9) to Eq.(3.8), we decompose the topical distribution of a review into the topical representations of the corresponding user and item:

$$\begin{aligned} \theta_{d,k} = P(k|u, i) &= P(k|u)P(k|i)\frac{P(u)P(i)}{P(k)P(u, i)} \\ &\propto P(k|u)P(k|i) = \theta_{u,k}\theta_{i,k} \end{aligned} \tag{3.10}$$

where $P(u), P(i)$ and $P(u, i)$ are constants in the LDA procedure, and the latent topics z_k are identically independent from each other, giving us constant and equal valued $P(k)$'s over the K topics. As a result, we have the following conditional recommendation probability for LDA models:

$$\begin{aligned} P_{\text{LDA}}(i|u) = \sum_{k=1}^{K} P(k|u)P(i|k) &= \sum_{k=1}^{K} P(k|u)P(k|i)\frac{P(i)}{P(k)} \\ &\propto \sum_{k=1}^{K} P(k|u)P(k|i) = \sum_{k=1}^{K} \theta_{u,k}\theta_{i,k} = \theta_u \cdot \theta_i \end{aligned} \tag{3.11}$$

and this result conforms with Eq.(3.10) in that, the probability of recommending an item given a user is positively correlated to the sum of topic probabilities that a user may textually review on an item.

3.2 Bridging the Objective Functions According to the conditional item recommendation probabilities given a target user specified in Eq.(3.6) and Eq.(3.11) for LFM and LDA models, respectively, a favoured approach to bridge the two models to leverage the power of both ratings and reviews should coordinate their objective functions, so that a higher value of $\gamma_u \cdot \gamma_i$ also corresponds to a higher value in $\theta_u \cdot \theta_i$. More precisely, except for the monotonic property defined in Eq.(2.4) on a vector itself, the randomization function $f(\cdot)$ from γ to θ is also required to be monotonic for vector multiplications:

(3.12)
$$\gamma_1 \cdot \gamma_2 < \gamma_3 \cdot \gamma_4 \rightarrow f(\gamma_1) \cdot f(\gamma_2) < f(\gamma_3) \cdot f(\gamma_4),\ \forall \gamma_1, \gamma_2, \gamma_3, \gamma_4$$

In this way, the LFM and LDA components in a bridged objective function would not conflict with each other during the model learning process, because both of them increase/decease the recommendation probability $P(i|u)$ at the same time for each single iteration.

Previous work intuitionally assumes that a randomization function satisfying the vector-level monotonic property in Eq.(2.4) will also be monotonic on product-level as Eq.(3.12). Frequently used examples are the normalization-based randomization functions, which normalize the elements of γ to construct θ so that they sum to one [19, 20, 22, 23]. In [19] for example, a logistic normalization randomization function as in Eq.(2.5) is applied so as to minimize the following joint objective function to bridge the LFM and LDA models:

$$\mathcal{O} = \sum_{r_{u,i} \in \mathcal{R}} \underbrace{\big(rate(u,i) - r_{u,i}\big)^2}_{\text{LFM component}} - \lambda \underbrace{\mathcal{L}\big(\mathcal{D}|\theta, \phi, z\big)}_{\text{LDA component}} \quad (3.13)$$

where the LFM component still minimizes the error in predicted ratings, while the LDA component is the log-likelihood of the probability for the review corpus in Eq.(2.3).

Previous designations do seem intuitional and reasonable, and they indeed improve the performance of personalized recommendation in many cases. However, we would like to point out in this work that the vector-level monotonic property does not necessarily guarantee the monotonic property on a product-level. Actually, we prove that such a randomization function that satisfies both vector- and product-level monotonic properties does not exist at all, and the proof is exposited in the Appendix of this paper.

For this reason, forcing a vector-level randomization function on the latent factors to bridge the LFM and LDA models will result in a conflict between the two components during the procedure of objective optimization, i.e., while the LFM component gains a higher probability of item recommendation with a larger value of $\gamma_u \cdot \gamma_i$, the LDA component may reversely force a lower recommendation probability with $\theta_u \cdot \theta_i$ just because of the mathematical property of the randomization function, which is not favoured in model learning process. This further explains the observation that the prediction accuracy of Eq.(3.13) tends to fluctuate drastically during optimization, although the overall performance generally tends to increase along with the iterations.

3.3 Direct Product-Level Randomization Despite that a randomization function with product-level monotonic property does not exist, we shall notice a simple fact that the de facto components that we need to consider so as to preserve the orderings of $P_{\text{LFM}}(i|u)$ and $P_{\text{LDA}}(i|u)$, are the final product of the latent factors or latent topics as a whole, i.e., $\gamma_u \cdot \gamma_i$ and $\theta_u \cdot \theta_i$, rather than each latent factor γ to a latent topic distribution θ separately.

More precisely, what we really need in the LDA model of Eq.(2.3) is the topic distribution of each document θ_d, where we have $\theta_{d,k} \propto \theta_{u,k}\theta_{i,k}$ by Eq.(3.10). As a result, we can apply a randomization function $f(\cdot)$ to the product of latent factors $\gamma_{u,k}\gamma_{i,k}$ directly, so as to obtain the product of latent topic distributions $\theta_{u,k}\theta_{i,k}$ as a whole, which is further positively correlated to $\theta_{d,k}$ that will finally be adopted by the LDA component for model learning.

A lot of normalization-based randomization functions guarantee the product-level monotonic property when applied to the product of latent factors directly. In this work, we adopt the logistic-normalization function to enforce $\theta_{u,k}\theta_{i,k}$ (and thus $\theta_{d,k}$) to be positive and sum to one:

(3.14)
$$\theta_{d,k} \propto \theta_{u,k}\theta_{i,k} = f(\gamma_{u,k}\gamma_{i,k}) \doteq \frac{\exp(\gamma_{u,k}\gamma_{i,k})}{\sum_{k'} \exp(\gamma_{u,k'}\gamma_{i,k'})}$$

which preserves the orderings of the dimensions from $\gamma_u \cdot \gamma_i$ to $\theta_u \cdot \theta_i$, and thus guarantees the positive correlation between $P_{\text{LFM}}(i|u)$ and $P_{\text{LDA}}(i|u)$ according to Eq.(3.6) and Eq.(3.11). Based on this direct product-level randomization, we are fortunately able to bridge the LFM and LDA models to leverage the power of ratings and reviews collectively, and meanwhile make the two components coordinate with each other for model learning, which improves both the performance and stability of personalized recommendation.

In the following, we describe our Joint Factoriza-

tional Topic (JFT) model, as well as its application in the practical scenario of (cross-city) restaurant recommendation.

4 Joint Factorizational Topic (JFT) Model

4.1 Model Design The JFT model bridges the LFM component as Eq.(2.1) and the LDA component in Eq.(2.3) according to the product-level randomization. Specifically, let $\theta_{d,k} = \frac{\exp(\gamma_{u,k}\gamma_{i,k})}{\sum_{k'}\exp(\gamma_{u,k'}\gamma_{i,k'})}$ in Eq.(2.3), the JFT model attempts to optimize the following objective function:

(4.15)
$$F(\Theta, \phi, z) = \sum_{r_{u,i}\in\mathcal{R}} \big(rate(u,i) - r_{u,i}\big)^2 - \lambda_l \mathcal{L}(\mathcal{D}|\theta, \phi, z) + \lambda_p \Omega(\Theta)$$

where $\Theta = (\alpha, \beta_u, \beta_i, \gamma_u, \gamma_i)$ is the parameter set of the LFM component, $\mathcal{L}(\mathcal{D}|\theta, \phi, z)$ is the log-likelihood of the whole corpus whose document distributions θ come from the latent factors γ, and $\Omega(\Theta) = \sum_{u,i}(\beta_u^2 + \beta_i^2 + \|\gamma_u\|_2^2 + \|\gamma_i\|_2^2)$ is the ℓ_2-norm regularizer for the latent parameters.

Intuitionally, the LDA component $\mathcal{L}(\cdot)$ serves as another regularization term besides the traditional ℓ_2-norm regularizer $\Omega(\cdot)$ for numerical rating prediction, and we trade off between them two with λ_l and λ_p, respectively. In this way, the JFT model attempts to minimize the error in rating prediction, and meanwhile maximizes the likelihood of observing the corresponding textual reviews.

Most importantly, the LFM and LDA components are designed to be consistent with each other by product-level randomization in model learning, in that a smaller prediction error in the LFM component functionally invokes a larger likelihood of observing the corresponding textual review, which makes the two components collaborate rather than violate with each other when optimizing the objective function.

4.2 Fitting the Model We introduce the algorithm for model fitting in this subsection. Typically, the LFM component (with ℓ_2-regularizer) can be easily fit with gradient descent, while the log-likelihood LDA component is optimized with Gibbs sampling. As our model jointly includes the two heterogenous components, we construct a learning procedure that optimizes the two components alternatively:

$S1$: $\{\Theta^t, \phi^t\} \leftarrow \underset{\Theta,\phi}{argmin}\, F(\Theta, \phi, z^{t-1})$ *by gradient descent*

$S2$: *Logistic normalization on each topic vector* ϕ_k^t

$S3$: *Sample* $z_{d,j}^t$ *with probability* $P(z_{d,j}^t = k) = \theta_{d,k}^t \phi_{k,w_{d,j}}^t$

In the first step, we fix the topic z of each word in each document, and further compute the gradient of each parameter in $\{\Theta, \phi\}$ while fixing the others. Based on these gradients, the parameters in $\{\Theta, \phi\}$ are updated one by one, where the step size for each parameter is determined by linear search.

However, the gradient descent procedure would not guarantee the word distribution ϕ of latent topics to be stochastic vectors. As a result, we conduct logistic normalization for each topic ϕ_k $(1 \le k \le K)$ in the second step, where each dimension of ϕ_k is normalized as $\phi_{k,w} = \frac{\exp(\phi_{k,w})}{\sum_{w'}\exp(\phi_{k,w'})}$.

In the last step, we preserve the results from the previous steps, and update the topic assignment for each word in each document. Similar to LDA, which assigns each word to the k-th topic according to the likelihood of the word discussing topic k, we set $z_{d,j} = k$ with probability proportional to $\theta_{d,k}\phi_{k,w_{d,j}}$, where the indices pair $\{d, j\}$ denotes the j-th word of document d, θ_d is the the topic distribution of document d, and ϕ_k is the word distribution of topic k.

The major difference between LDA and the last step of our JFT model is that, the topic distributions θ_d are determined base on the product-level randomization from latent factors γ_u, γ_i and γ_c in our model, instead of sampling from a Dirichlet distribution in LDA. As a result, we only need to sample the topic assignments z in each iteration of our JFT model. The probabilistic interpretation of our approach and its relationship with the LFM component are exposited in previous sections.

Finally, these steps are repeated iteratively until convergence, i.e., the ℓ_2-difference in Θ is sufficiently small between two consecutive iterations, or that an overfitting is observed in the validation set.

4.3 Top-N Recommendation We further adapt the JFT model to provide practical top-N recommendation lists beyond numerical rating predictions.

It is known that a good performance on rating prediction does not necessarily guarantee a satisfactory performance of top-N recommendation by ranking the items in descending order of the predicted ratings [27]. This is partly because of the contradiction between the goal of recommending items that users would potentially visit and the data (ratings) that we use for model training, i.e., users indeed visited the items in the dataset, no matter what numerical ratings they eventually made on them. Intuitionally, a relatively low predicted rating does not necessarily mean that the user would not be attracted by the item at all, because of the many items with low ratings yet visited by the users.

As a result, we train our JFT model (and the baseline approaches) for top-N recommendation in a differ-

ent way from the rating prediction task. Specifically, we feed the learning procedure with binary inputs, where the observed records in $\mathcal{R}$ are treated as positive cases (rating=1), and the negative cases (rating=0) are sampled from the unobserved user-item pairs in a balanced negative sampling manner. For clarity, we exposit the sampling, learning, and recommendation produce in Algorithm 1.

Algorithm 1 Top-N Recommendation

Input: $\mathcal{R}, \mathcal{D}, N$
Output: Recommendation list of length N
$\mathcal{R}^+ \leftarrow \mathcal{R}$ with all ratings reset to be 1
Initialize model parameters $\alpha, \beta, \gamma, \phi, z$ randomly
while Convergence **or** $t > T$ **do**
  $t \leftarrow t + 1$, $\mathcal{R}^- \leftarrow \emptyset$
  for $(u, i) \in \mathcal{R}^+$ **do**
    Sample item j from the same city of item i randomly, where $(u, j) \notin \mathcal{R}^+$
    $\mathcal{R}^- \leftarrow \mathcal{R}^- \cup (u, j)$ with rating 0
  end for
  Update model by $Step1 \sim 3$ with $\{\mathcal{R}^+ \cup \mathcal{R}^-, \mathcal{D}\}$
end while
Rank the items in descending order of predictions
return Top-N Recommended items for each user

5 Experiments

In this section, we conduct extensive experiments to validate the performance of our JFT model based on real-world datasets. We first conduct case studies to support the underlying intuition of our approach in cross-city recommendation, and further investigate the performance on rating prediction and top-N recommendation

5.1 Experimental Setup We collected user reviews from a major restaurant review website Dianping.com in China, including 253,749 reviews from 32,529 users towards 8,026 restaurants located in 194 cities, where each user made 20 or more reviews, including intro- and cross-city cases, and the home city of a user is registered in his/her profile. Of the whole 253,749 reviews, 233,802 (92%) records fall into intro-city reviews, and the remaining 19,947 (8%) records are cross-city reviews. Each review in the corpus consists of an integer rating ranging from 1 to 5 stars and a piece of textual comment, where the user expresses his/her opinions on the corresponding restaurant. The average length of textual comments is 41.5 words. To feed the LDA component with high quality textual inputs, we conduct part-of-speech tagging and stop word removing for each review with the widely used Stanford NLP toolkit.

After careful tuning with grid search, we set the hyper-parameters $\lambda_l = 0.01$ and $\lambda_p = 0.001$, and five-fold cross-validation was conducted in performance evaluation for all methods.

5.2 Case Study of Regional Features Before introducing the experimental results, we execute an intuitional case study of the distribution of regional features in support of the underlying intuition of cross-city recommendation based on textual reviews.

To do so, we select four representative feature words from restaurant reviews (*Mutton, Seafood, Spicy*, and *Price*) [28, 29], and calculate the percentage that reviews in each city discussing a feature against the total number of reviews discussing the feature in the whole corpus, which represent the regional distributions of a feature among the major cities across mainland China. We further intuitionally plot the regional distributions for each feature on maps, where each dot represents a city, and the grey-scale color is proportional to its percentage of discussing a feature [30], as shown in Figure 1.

It is observed that the feature *Mutton* is mainly discussed in cities located in the north and west part of China, which is consistent with the distribution of the grasslands in China. *Seafood* is mostly discussed along the seacoast, while *Spicy* is extensively mentioned in provinces that are famous for spicy food, including Sichuan, Hunan, Guangxi, Shaanxi, Xinjiang, etc. Finally, the feature *Price* is concerned by customers in nearly all the cities, which is also reasonable for its general importance to most users, though they have different flavour preferences.

We further display the top 10 words of each topic discovered by the LDA component with $K = 10$, as shown in Table 3, and find that they exhibit clear meanings in different Cuisine styles (Cantonese, Northwest, Northeast, Western, etc.) or food types (Seafood, Meat, Drinks, Desserts, etc.).

5.3 Performance on Rating Prediction In this section, we study the performance on rating prediction tasks, and adopt the following state-of-the-art rating prediction baselines.

LFM: The LFM approach denoted in Eq.(2.2), which takes no advantage of the textual reviews.

EFM: The Explicit Factor Model presented in [15], which is the state-of-the-art recommendation approach based on textual reviews by sentiment analysis.

HFT: The Hidden Factors and Topics model in [19], which is the state-of-the-art method that takes advantages of both LFM and LDA by applying vector-level randomizations (Eq.(2.5)).

We adopt Root Mean Square Error (RMSE) and

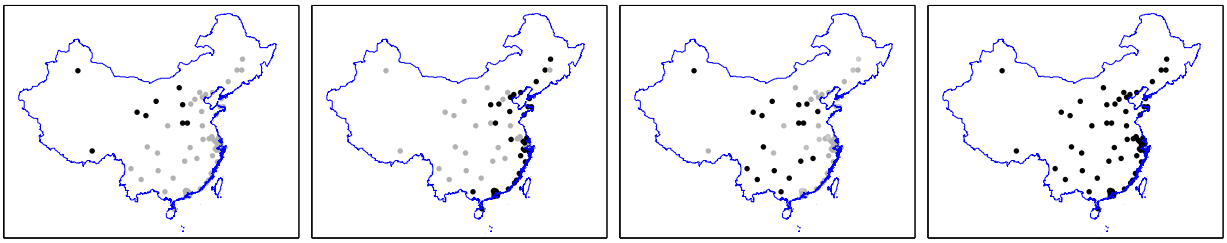

(a) Mutton (b) Seafood (c) Spicy (d) Price

Figure 1: Location distribution of some representative feature words. Each dot on the map represents a city, and its grey-scale color is proportional to the frequency that the word is discussed in the reviews of that city.

Mean Absolute Error (MAE) for evaluation, and the results with the number of topics/factors $K = 10$ are shown in Table 2. The standard deviations in five-fold cross-validation for each method and metric are ≤ 0.002.

Table 2: RMSE and MAE when $K = 10$. Standard deviations for each method are ≤ 0.002.

Method	LFM	EFM	HFT	JFT
RMSE	0.6688	0.6529	0.6532	**0.6386**
MAE	0.5309	0.5283	0.5280	**0.5128**

We find that all the other approaches gain better performance against LFM, which means that taking advantage of the textual reviews helps to make better rating predictions. Beyond this, our JFT model achieves the best performance against both EFM and HFT. On considering that a major difference between JFT and HFT is the product- and vector-level randomization, this result is consistent with the theoretical analysis to bridge the LFM and LDA components in Section 3.

To exhibit a clearer view of the performance on cross-city scenarios, we further take out the cross-city rating records from the test set in each of the 5 folds, and conduct performance evaluation under different choices of the number of latent factors and topics K from 10 through 100. Results for RMSE and MAE are shown in Figure 2. We see that our JFT approach outperforms all other baselines on all choices of topic numbers, which validates the superior performance when we consider the local features of a city from reviews and match them to latent factors in consistent manners.

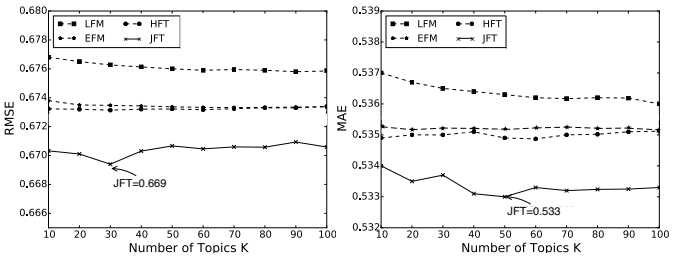

(a) RMSE vs Topic Number (b) MAE vs Topic Number

Figure 2: RMSE and MAE vs the number of topics or latent factors K in cross-city settings.

5.4 Top-N Recommendation In this subsection, we explore the performance of our approach in more practical top-N recommendation tasks, and make comparison with the following baseline methods:

WRMF: Weighted Regularized Matrix Factorization [31], which is similar to LFM but applies weighted negative sampling to benefit top-N recommendations.

BPRMF: Bayesian Personalized Ranking (BPR) for MF [32], which is the state-of-the-art algorithm for top-N recommendation based only on numerical ratings.

HFT: The original HFT method does not achieve satisfactory top-N performance in our settings. As a result, we improve it with the same binary inputs and negative sampling approach (Alg.1) for fair comparison.

To evaluate, we randomly hold out 5 records for each user, and provide top-5 recommendation list for each user, as with most practical applications. We adopt the measures of Precision@5 and NDCG@5, where the latter takes the positions of recommended items into consideration, and the results are shown in Table 4.

Table 4: Prec@5 and NDCG@5 with $K = 10$. Standard deviations for each method are ≤ 0.0006.

Method	WRMF	BPRMF	HFT	JFT
Precision@5(%)	0.60	0.57	0.65	**0.79**
NDCG@5	0.1616	0.1632	0.1653	**0.1790**

We see that both our JFT approach and the HFT method (which make use of textual reviews) gain better performance than WRMF and BPRMF (which only make use of ratings). Further more, our JFT method gains a 22% improvement against HFT in terms of precision, and 8.3% on NDCG, which is a superior achievement for practical applications.

Similar to the task of rating prediction, we also

Table 3: Top ten words of each topic discovered by the LDA component when the number of topics $K = 10$. Words in the first row are manually assigned labels to summarize the corresponding topic.

General	Dessert	Spicy	Seafood	Cantonese	Northwest	Northeast	Meat	Drinks	Western
waiters	milk tea	hotpot	seafood	price	mutton	cold noodle	steak	beer	salad
taste	cake	Chongqing	seaside	breakfast	soul milk	roast duck	taste	corn	service
service	salad	discount	pie	seafood	donkey	roast meat	business	taro	buffet
famous	cheese	caffe	mackerel	tofu	muslim	dumpling	intestine	drink	steak
environment	bread	spicy	globefish	food	cold dish	free	ribs	tea	pizza
business	dumpling	vegetables	crab	HongKong	hotpot	potato	chicken	environment	seafood
price	mango	rice	scallop	tea	noodle	flavour	mutton	peanut butter	feeling
discount	chocolates	drink	lobster	GuangZhou	Tibet	cabbage	durion	melon seed	sushi
feeling	dessert	Guizhou	holothurian	cheese	rice noodle	Beijing	apple	milk tea	burger
favour	cream	chicken wings	shrimp	roast goose	pilau	ribs	pie	green tea	food

evaluate the top-N performance in cross-city settings. To do so, we select those user-city pairs that a user has at least 5 records in a city beyond his/her home city, which results into 4,021 pairs corresponding to 1,739 users. We thus randomly hold out 5 records for a user in a corresponding city, and construct the recommendation list using the restaurants from that city, which gives us 4,021 lists for evaluation. We also conduct 5-fold cross-validation, and the standard deviations for both Precision and NDCG are ≤ 0.005. Figure 3 shows the results against the number of latent factors/topics K.

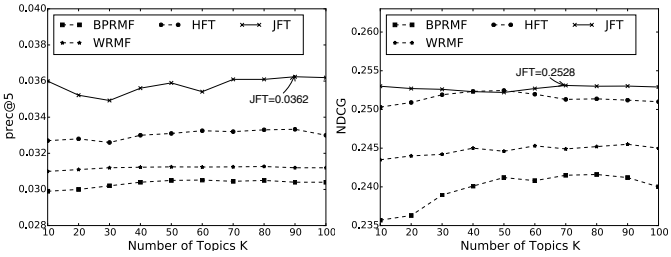

(a) Precision vs Topic Number (b) NDCG vs Topic Number

Figure 3: Precision@5 and NDCG@5 vs the number of topics or latent factors K in cross-city settings.

We see that the performance of our JFT model is better than the baselines on nearly all choices of K. Interestingly, we find that the overall cross-city performance is a magnitude better than that on the whole dataset. This means that user behaviours can be more predictable in cross-city settings, where users do visit local attractions beyond their historical preferences.

6 Related Work

With the continuous growth of various online items across a vast range of the Web, Personalized Recommender Systems (PRS) [1] have set their missions to save users from information overload [2], and they have been widely integrated into various online applications in the forms of, for example, product recommendation in e-commerce [33], friend recommendation in social networks [34], news article recommendation in web portals [35], and video recommendation in video sharing websites [36], etc.

Early systems of personalized recommendations rely on content-based approaches [37], which construct the user/item content profiles and make recommendation by paring users with the contently similar items. Content-based approaches usually gain good accuracy but functionally lack the ability of providing recommendations with novelty, serendipity, and flexibility. Besides, they usually require a large amount of expensive human annotations [1].

This further leads to the prospering of Collaborative Filtering (CF)-based recommendation algorithms [4, 3] that leverage the wisdom of the crowds. Typically, they construct the partially observed user-item rating matrix and conduct missing rating prediction based on the historical records of a user, as well as those of the others.

With widely recognized performance in rating prediction, scalability, and computational efficiency, the Latent Factor Models (LFM) [5] based on Matrix Factorization (MF) [6, 7] techniques for CF have been extensively investigated by the research community, and widely applied in practical systems. Perhaps the most early and representative formalization of LFM for recommendation dates back to Koren et al [26], and other variants for personalization include Non-negative Matrix Factorization (NMF) [8], Probabilistic Matrix Factorization (PMF) [9, 38], and Maximum Margin Matrix Factorization (MMMF) [39, 40], etc.

Despite the important success in rating prediction, the CF approaches based solely on the numerical ratings suffer from the problems of explainability [15], cold-start [41, 42], and the difficulty to provide more specific recommendations that meet targeted item aspects [43].

Besides, related research results show that the performance on numerical rating prediction does not necessarily relate to the performance on practical top-N recommendations [27], and that the numerical star ratings may not always be a reliable indicator of users' attitudes towards items [44].

To alleviate these problems, researchers have been investigating the incorporation of textual reviews for recommendation [45, 46, 47, 48], which is another important information source beyond the star ratings in many systems. Early approaches rely on manually extracted item aspects from reviews for more informed recommendation [49, 43] and rating prediction [50, 51], which improved the performance but also required extensive human participations. As a results, researchers recently have begun to investigate the possibility of integrating the automatic topic modeling techniques on textual reviews and the latent factor modeling approach on numerical ratings for boosted recommendation, and have achieved appealing results [19, 20, 22, 23, 52].

However, without a clear mathematical exposition of the relationships between latent factor models and topic modeling, current approaches have to base themselves on manually designed randomization functions or probabilistic distributions. In this work, however, we attempt to make an exposition on the relationships between the two types of objective functions, and further bridge the inherently heterogenous models in a harmonious way for recommendation with the power of both numerical ratings and textual reviews.

7 Conclusions and Future Work

In this paper, we proposed the Joint Factorizational Topic (JFT) model that leverages the ratings and reviews in a collective manner for recommendation. For the first time, we examine the mathematical mutual-dependent relations between LFM and LDA for personalized recommendation, and we prove that vector-level randomization functions that are multiplicatively monotonic actually do not exist, although they are previously frequently used to bridge the heterogeneous LFM and LDA components.

Fortunately, we also find that a direct product-level randomization approach can be used to bridge the two components and coordinate their behaviors for model learning. Extensive experimental results on case studies, rating prediction, and top-N recommendation tasks verified both the intuition, theoretical basis, and the quantitative performance of our approach.

This is the first step towards bridging the latent factor models and topical models for cross-city recommendation, and there is much room for future research and improvements. One important thing to notice is that our model is not restricted to cross-'city' scenarios, but generally applicable to cross-domain recommendations, and the appealing performance of our bridging framework sheds light on possibly new approaches for the traditional task of cross-domain recommendation. Besides the factorizational and topic models, we can even investigate other graphical or deep-representational models to integrate the numerical ratings and topics from textual reviews, so as to bridge the two important personalized information sources broadly seen in many Web-based systems.

Appendix

Let γ denote arbitrary vectors with length K, and $\mathcal{F}$ is the set of all randomization functions $f : \mathbb{R}^K \to \mathbb{R}^K$ satisfying:

(7.16)
$$\begin{cases} 0 \le f(\gamma)_i \le 1, \|f(\gamma)\|_1 = 1 \\ \gamma_i < \gamma_j \to f(\gamma)_i < f(\gamma)_j \end{cases}, \forall \gamma \in \mathbb{R}^K, 1 \le i, j \le K$$

then there exists no randomization function $f \in \mathcal{F}$ with the product-level monotonic property of:

(7.17)
$$\gamma_1 \cdot \gamma_2 < \gamma_3 \cdot \gamma_4 \to f(\gamma_1) \cdot f(\gamma_2) < f(\gamma_3) \cdot f(\gamma_4),\ \forall \gamma_1, \gamma_2, \gamma_3, \gamma_4$$

Proof: Suppose there exists a randomization function $f \in \mathcal{F}$ that meets Eq.(7.17). Let $t > 1$, and let α and β be vectors with $\alpha \cdot \beta > 0$, then we have $t\alpha \cdot \beta > \alpha \cdot \beta$. By applying the property of product-level monotonic in Eq.(7.17) we have:

$$f(t\alpha) \cdot f(\beta) > f(\alpha) \cdot f(\beta) \tag{7.18}$$

and this can be equivalently written as:

$$\big(f(t\alpha) - f(\alpha)\big) \cdot f(\beta) > 0 \tag{7.19}$$

Let $\Delta \doteq f(t\alpha) - f(\alpha)$, and according to the definition of randomization function in Eq.(7.16), we know that $\sum_k f(t\alpha)_k = \sum_k f(\alpha)_k = 1$, thus we have:

$$\sum\nolimits_k \Delta_k = \sum\nolimits_k \big(f(t\alpha) - f(\alpha)\big)_k = 0 \tag{7.20}$$

According to Eq.(7.19) we know that $\Delta \neq \mathbf{0}$. Let $\mathcal{P}$ denote the indices of all positive elements in vector Δ, and $\mathcal{N}$ be the indices of negative elements. We have:

$$\sum\nolimits_{k \in \mathcal{P}} \Delta_k + \sum\nolimits_{k \in \mathcal{N}} \Delta_k = 0 \tag{7.21}$$

As Eq.(7.19) holds for any β with $\alpha \cdot \beta > 0$, without loss of generally, let β be a vector where $\beta_{k \in \mathcal{P}} = 0$ and $\beta_{k \in \mathcal{N}} = 1$. According to the vector-level monotonic property in Eq.(7.16) and the fact that $0 < 1$, we have $f(\beta)_{k \in \mathcal{P}} < f(\beta)_{k \in \mathcal{N}}$ and $0 \le f(\beta)_{k \in \mathcal{P} \cup \mathcal{N}} \le$

1. Combined with Eq.(7.21), we further obtain the following:

$$\Delta \cdot f(\beta) = \sum_{k \in \mathcal{P}} \Delta_k f(\beta)_k + \sum_{k \in \mathcal{N}} \Delta_k f(\beta)_k < 0 \tag{7.22}$$

which is a direct contradiction with Eq.(7.19). As a result, there exists no randomization function $f \in \mathcal{F}$ that satisfies the product-level monotonic property in Eq.(7.17). □

References

[1] F. Ricci, L. Rokach, and B. Shapira, "Introduction to Recommender Systems Handbook," *Springer*, 2011.

[2] J. A. Konstan and J. Riedl, "Recommender Systems: from Algorithms to User Experience," *User Modeling and User-Adapted Interaction*, vol. 22, no. 1-2, pp. 101–123, 2012.

[3] X. Su and T. M. Khoshgoftaar, "A Survey of Collaborative Filtering Techniques," *Advances in AI*, vol. 4, 2009.

[4] Y. Koren and R. Bell, "Advances in Collaborative Filtering," *Recommender Systems Handbook*, pp. 145–186, 2011.

[5] M. Knott and D. Bartholomew, "Latent Variable Models and Factor Analysis," *Kendall's Library of Statistics*, vol. 2, 1999.

[6] G. Takacs, I. Pilaszy, B. Nemeth, and D. Tikk, "Investigation of Various Matrix Factorization Methods for Large Recommender Systems," *Proc. ICDM*, 2008.

[7] A. P. Singh and G. J. Gordon, "A Unified View of Matrix Factorization Models," *Machine Learning and Knowledge Discovery in Databases*, pp. 358–373, 2008.

[8] D. D. Lee and H. S. Seung, "Algorithms for Non-negative Matrix Factorization," *Proc. NIPS*, 2001.

[9] R. Salakhutdinov and A. Mnih, "Probabilistic Matrix Factorization," *Proc. NIPS*, 2008.

[10] Y. Zhang, M. Zhang, Y. Liu, S. Ma, and S. Feng, "Localized Matrix Factorization for Recommendation based on Matrix Block Diagonal Forms," *WWW*, 2013.

[11] Y. Zhang, M. Zhang, Y. Liu, and S. Ma, "Improve Collaborative Filtering Through Bordered Block Diagonal Form Matrices," *SIGIR*, 2013.

[12] Y. Zhang, M. Zhang, Y. Zhang, Y. Liu, and S. Ma, "Understanding the Sparsity: Augmented Matrix Factorization with Sampled Constraints on Unobservables," *CIKM*, 2014.

[13] Y. Zhang, M. Zhang, Y. Liu, C. Tat-Seng, Y. Zhang, and S. Ma, "Task-based recommendation on a web-scale," in *Big Data (Big Data), 2015 IEEE International Conference on*. IEEE, 2015, pp. 827–836.

[14] Y. Zhang, "Browser-oriented universal cross-site recommendation and explanation based on user browsing logs," in *Proceedings of the 8th ACM Conference on Recommender systems*. ACM, 2014, pp. 433–436.

[15] Y. Zhang, G. Lai, M. Zhang, Y. Zhang, Y. Liu, and S. Ma, "Explicit Factor Models for Explainable Recommendation based on Phrase-level Sentiment Analysis," *SIGIR*, 2014.

[16] X. Xu, A. Datta, and K. Dutta, "Using Adjective Features from User Reviews to Generate Higher Quality and Explainable Recommendations," *IFIP Advances in Info. and Com. Tech.*, vol. 389, pp. 18–34, 2012.

[17] D. M. Blei, "Probabilistic Topic Models," *Communications of the ACM*, vol. 55, no. 4, pp. 77–84, 2012.

[18] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent Dirichlet Allocation," *JMLR*, vol. 2003, no. 3, pp. 993–1022, 2003.

[19] J. McAuley and J. Leskovec, "Hidden Factors and Hidden Topics: Understanding Rating Dimensions with Review Text," *RecSys*, pp. 165–172, 2013.

[20] G. Ling, M. R. Lyu, and I. King, "Ratings meet Reviews: a Combined Approach to Recommend," *RecSys*, 2014.

[21] X. Chen, P. Wang, Z. Qin, and Y. Zhang, "Hlbpr: A hybrid local bayesian personal ranking method," in *Proceedings of the 25th International Conference Companion on World Wide Web*. International World Wide Web Conferences Steering Committee, 2016, pp. 21–22.

[22] D. Agarwal and B.-C. Chen, "fLDA: Matrix Factorization through Latent Dirichlet Allocation," *WSDM*, 2010.

[23] C. Wang and D. M. Blei, "Collaborative Topic Modeling for Recommending Scientific Articles," *KDD*, 2011.

[24] S. Purushotham, Y. Liu, and C.-C. J. Kuo, "Collaborative Topic Regression with Social Matrix Factorization for Recommendation Systems," *ICML*, 2012.

[25] X. Chen, Z. Qin, Y. Zhang, and T. Xu, "Learning to rank features for recommendation over multiple categories," in *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*. ACM, 2016, pp. 305–314.

[26] Y. Koren, R. Bell, and C. Volinsky, "Matrix Factorization Techniques for Recommender Systems," *Computer*, 2009.

[27] P. Cremonesi, Y. Koren, and R. Turrin, "Performance of Recommender Algorithms on Top-N Recommendation Tasks," *RecSys*, pp. 39–46, 2010.

[28] Y. Tan, Y. Zhang, M. Zhang, Y. Liu, and S. Ma, "A Unified Framework for Emotional Elements Extraction based on Finite State Matching Machine," *NLPCC*, vol. 400, pp. 60–71, 2013.

[29] Y. Zhang, M. Zhang, Y. Liu, and S. Ma, "Boost phrase-level polarity labelling with review-level sentiment classification," *arXiv preprint arXiv:1502.03322*, 2015.

[30] H. Zhao, P. Poupart, Y. Zhang, and M. Lysy, "Sof: Soft-cluster matrix factorization for probabilistic clustering." in *AAAI*, 2015.

[31] Y. Hu, Y. Koren, and C. Volinsky, "Collaborative Filtering for Implicit Feedback Datasets," *ICDM*, 2008.

[32] S. Rendle, C. Freudenthaler, Z. Gantner, and L. S. Thieme, "BPR: Bayesian Personalized Ranking from Implicit Feedback," *UAI*, 2009.

[33] G. Linden, B. Smith, and J. York, "Amazon.com Recommendations: Item-to-item Collaborative Filtering," *Internet Computing*, vol. 7, no. 1, pp. 76–80, 2003.

[34] E. A. Baatarjav, S. Phithakkitnukoon, and R. Dantu, "Group Recommendation System for Facebook," *On the Move to Meaningful Internet Systems: OTM 2008 Workshops*, pp. 211–219, 2008.

[35] A. Das, M. Datar, A. Garg, and S. Rajaram, "Google News Personalization: Scalable Online Collaborative Filtering," *WWW*, pp. 271–280, 2007.

[36] J. Davidson, B. Liebald, J. Liu *et al.*, "The YouTube Video Recommendation System," *RecSys*, pp. 293–296, 2010.

[37] M. Pazzani and D. Billsus, "Content-Based Recommendation Systems," *The Adaptive Web LNCS*, pp. 325–341, 2007.

[38] R. Salakhutdinov and A. Mnih, "Bayesian Probabilistic Matrix Factorization using Markov Chain Monte Carlo," *Proc. ICML*, 2008.

[39] N. Srebro, J. D. M. Rennie, and T. S. Jaakkola, "Maximum-Margin Matrix Factorization," *NIPS*, 2005.

[40] J. Rennie and N. Srebro, "Fast Maximum Margin Matrix Factorization for Collaborative Prediction," *ICML*, 2005.

[41] B. Lika, K. Kolomvatsos, and S. Hadjiefthymiades, "Facing the Cold Start Problem in Recommender Systems," *Expert Systems with Applications*, vol. 41, no. 4, pp. 2065–2073, 2014.

[42] A. Levi, O. Mokryn, C. Diot, and N. Taft, "Finding a Needle in a Haystack of Reviews: Cold Start Context-Based Recommender System," *RecSys*, pp. 115–122, 2012.

[43] M. Ko, H. W. Kim, M. Y. Yi, J. Song, and Y. Liu, "MovieCommenter: Aspect-Based Collaborative Filtering by Utilizing User Comments," *CollaborateCom*, 2011.

[44] Y. Zhang, H. Zhang, M. Zhang, Y. Liu *et al.*, "Do Users Rate or Review? Boost Phrase-level Sentiment Labeling with Review-level Sentiment Classification," *SIGIR*, 2014.

[45] Y. Zhang, "Incorporating phrase-level sentiment analysis on textual reviews for personalized recommendation," in *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*. ACM, 2015, pp. 435–440.

[46] M. Terzi, M. A. Ferrario, and J. Whittle, "Free text in user reviews: Their role in recommender systems," *RecSys*, 2011.

[47] Y. Zhang, H. Zhang, M. Zhang, Y. Liu, and S. Ma, "Do users rate or review?: boost phrase-level sentiment labeling with review-level sentiment classification," in *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*. ACM, 2014, pp. 1027–1030.

[48] Y. Zhang, M. Zhang, Y. Zhang, G. Lai, Y. Liu, H. Zhang, and S. Ma, "Daily-aware personalized recommendation based on feature-level time series analysis," in *Proceedings of the 24th International Conference on World Wide Web*. ACM, 2015, pp. 1373–1383.

[49] S. Aciar, D. Zhang, S. Simoff, and J. Debenham, "Informed Recommender: Basing Recommendations on Consumer Product Reviews," *Intelligent Systems*, vol. 22, no. 3, pp. 39–47, 2007.

[50] N. Jakob, S. H. Weber, M. C. Mller *et al.*, "Beyond the Stars: Exploiting Free-Text User Reviews to Improve the Accuracy of Movie Recommendations," *TSA*, 2009.

[51] G. Ganu, N. Elhadad, and A. Marian, "Beyond the Stars: Improving Rating Predictions using Review Text Content," *WebDB*, 2009.

[52] Y. Zhang, Y. Tan, M. Zhang, Y. Liu, C. Tat-Seng, and S. Ma, "Catch the black sheep: unified framework for shilling attack detection based on fraudulent action propagation," in *Proceedings of the 24th International Conference on Artificial Intelligence*. AAAI Press, 2015, pp. 2408–2414.